

10/528598

Rec. TO 21 MAR 2005

Method for routing data packets, and devices for implementing the method

5 The invention concerns a method for routing data frames, for example Ethernet frames, in an Ethernet routing device. The invention applies to DSL modems, but is not limited to that environment.

10 IP multicasting is the transmission of an IP datagram to a group of devices identified by a single IP destination address. A device connected to a LAN (Local Area Network), for example an Ethernet LAN, may join or leave a group, i.e. request that multicast packets concerning a certain group be addressed to it or not. The Internet Group Management Protocol ('IGMP') provides a way for a device to report its multicast group membership to adjacent
15 routers.

 The LAN can be connected to a WAN (Wide Area Network) through a routing device. This routing device, which may be an ADSL modem, has to learn which packets received on the WAN are to be forwarded on the LAN in the context of multicast groups. For this purpose, it can monitor IGMP packets
20 sent from the devices on the LAN to the WAN routers. This process, called 'snooping', allows the routing device to update its internal tables for the purpose of filtering packets from the WAN.

 The number of IGMP packets may be fairly important. If a switch module connected to the LAN communicates IGMP packets to the multicast
25 management control module of the routing device through control busses, the bandwidth limitations might be too important. On the other hand, the processing of frames – IGMP or not – should not be made more complex simply because of any special treatment of the IGMP frames.

30 The invention concerns a method for routing data packets in a routing device connecting a first network and a second network, characterized by the steps of:

- (a) receiving a frame from a device connected to the first network;
- (b) forwarding the frame to an internal bridge module of the device;
- 35 (c) checking whether the frame contains a multicast group management message and in the affirmative, creating a new frame comprising as destination address the destination address of an internal multicast group

management module and as payload at least the multicast management data of the received frame; and

(d) forwarding this new frame to the internal bridge function.

5 By creating a new frame and sending it through the internal bridge module in the same fashion as the received frames, there is no need to use any low bit-rate control bus: the normal data bus may be used. The internal bridge module does not have to carry out any specific processing either for delivering the information to the right module. All packets, whether concerning multicast
10 group management messages or not, are handled alike by the internal bridge function.

According to an embodiment of the invention, the first network is an Ethernet network and the steps (a) to (d) are carried out by an Ethernet switch
15 module.

According to an embodiment of the invention, the method further comprises the step of inserting into the new frame an identifier of a port on which the initial frame was received.
20

According to an embodiment of the invention, the multicast group management message is an IGMP message.

Another object of the invention is a routing device for connecting a first and a second network, characterized in that it comprises:
25

(a) a switch for receiving frames from the first network;

(b) an internal bridge function for forwarding frames to appropriate modules, said bridge function being connected to the switch;

(c) a multicast group management module for maintaining up to date
30 multicast group information based on frames received on the first network, said multicast group management module being connected to the bridge function for receiving selected frames there from;

wherein the switch is adapted to determine whether a received frame comprises a multicast group management message, and in the affirmative,
35 providing a new frame comprising multicast group management information extracted from the original frame, wherein the second frame has a destination

address equal to the address of the multicast group management module, and for delivering the new frame to the bridge function.

Another object of the invention is a switch circuit for use in a device as defined above, said circuit comprising multiple input ports, characterized in that it is adapted to determine whether a received frame comprises a multicast group management message, and in the affirmative, providing a new frame comprising multicast group management information extracted from the original frame, wherein the second frame has a destination address equal to the address of the multicast group management module, and for forwarding the new frame to the bridge function.

Other characteristics and advantages of the invention will appear in the description of the non-restricting embodiment, explained with the help of the enclosed figures, among which:

- figure 1 is a block diagram of the device according to the present embodiment;
- figure 2 is a diagram of an Ethernet frame according to the present embodiment;
- figure 3 is a flowchart of the method according to the present embodiment.

More information about multicasting can be found in the documents 'IETF RFC 1112' and 'IETF RFC 2236' (Internet Engineering Task Force Request for Comments).

Figure 1 is a schematic block diagram of a modem 1 connecting a Local Area Network (LAN) and a Wide Area Network (WAN).

The DSL modem comprises an Ethernet circuit 2 and an integrated chipset 3. The Ethernet switch 2 handles the Local Area Network (LAN) traffic received/sent through its ports. In the present embodiment, the LAN is an Ethernet network. The integrated chipset 3 contains multiple blocks: a bridge circuit 4 able to bridge incoming traffic to all bridge interfaces of the modem (namely here, ATM PVCs – for Asynchronous Transfer Mode Permanent Virtual Circuits – used in a 'bridged' mode), a 'packets and cells' engine 5 responsible for ATM and DSL encapsulation, transmission and reception and a CPU 6 that

runs software for specific processing needs. The Ethernet switch 2 and the IC 3 are connected through a data bus 7 and a management and control bus 8.

The Ethernet switch receives Ethernet frames from the devices connected to its ports, e.g. device 9, which is for example a personal computer.

5 Frames received by the switch are forwarded to the IC 3 through the data bus 7.

According to the present embodiment, the Ethernet switch detects IGMP IP frames in the Ethernet frames transmitted by devices (not shown) on the LAN. IGMP frames are IP datagrams which IP protocol number is '2' (IGMP). When such a frame is detected, the Ethernet switch delivers the
10 original frame to the IC 3, the Ethernet layer being removed in the CPU, when a packet is handled by the latter (and also on the destination device in the WAN). It further creates a new frame, incorporating the detected frame, as well as additional information. According to the present embodiment, this additional information is the port number of the switch on which the IGMP frame has been
15 received. The switch transmits both the additional port information as well as the IGMP frame to the CPU 6, through the data bus. For this purpose, the switch creates a new Ethernet frame, having a destination address equal to that of the CPU 3, and containing in its payload the port information, e.g. port number, and the IGMP Ethernet frame. An example of such a frame is shown in
20 figure 2. The frame is identified as being a forwarded IGMP packet with port information by a specific Type value.

The bridge circuit receives this frame, and given the destination address, transmits it to the CPU. It does not deliver this frame on the WAN. The CPU detects the presence of the port number and IGMP frame in the payload. It
25 then uses the contents of the IGMP frame, in particular the multicast destination address, to update its multicast group information for the devices on the LAN.

If the message is a 'join group' message and the group already exists, the port number of the incoming frame is added to the group members. The information is thus transmitted to the switch, which is a managed switch
30 having a feature that allows grouping of ports. Sending to the group is then equivalent to sending to each port being part of that group.

If the message is a 'join group' message and the group does not yet exist, it is created, and the port number of the incoming frame is added as a first group member.

35 If the message is a 'leave group' message, the source address is removed from the group.

The multicast information is maintained in a memory (not illustrated) by the CPU, under the form of appropriate tables.

The above process is illustrated by the flowchart of figure 3.

5

The CPU filters IP frames having as destination address one of the multicast addresses of the groups of which devices on the LAN are members, and to transmit filtered frames on the network to the appropriate ports. The packet/cell engine 5 transmits and receives packets respectively to and from the DSL line.

10

Although the embodiment of the invention concerns an Ethernet and IP environment, the invention is not limited to this particular case. Moreover, other messages than IGMP messages may be processed in a similar way, and any other multicast management protocol can be used. The invention can be used in other environments when a packet of data needs to be forwarded to one or more modules or processes, depending on packet content, without blocking the default path of the packet.

15

20